

**COMPROMISES ALONG THE WAY:
BALANCING SPEED TO MARKET WITH
SUSTAINABILITY WHILE DELIVERING KNOWLEDGE
MANAGEMENT SERVICES***

By Martha K Heyman

Abstract: Knowledge management in a consulting firm is where the rubber meets the road. It is where information science standards and guidelines are compromised in favor of business processes and requirements. The challenge is 'exasperated' by the use of a single, shared knowledge asset repository across the firm. The needs and work processes of each practice area within a consulting firm vary greatly. The common technology platform and repository require that compromises be reached between not just information science standards and day-to-day business operations, but also between the different knowledge managers who support the firm's different practice areas. This paper will discuss some of the compromises, and the path to those compromises, which must be made while implementing a successful knowledge management program within a for-profit enterprise. The paper will also briefly contrast the structure of a successful knowledge management program within a consulting firm versus one within a commodities business, specifically a global chemical manufacturing company.

Why Knowledge Management?

There is a temptation to assume business value for knowledge management exists just because it has become a business buzzword, as if the presence of a slick term makes business value a fore-gone conclusion. Remember the old adage from your mother: "Just because the other kids are doing it, does not mean you should"? Every organization, whether it is a for-profit enterprise or a not-for-profit organization, must take the time to think for itself. Is there a business case for us to manage our knowledge? That said, however, I would venture to say that if the conclusion from the self-assessment is no, then either the assessment or the thinking is flawed. Certainly no enterprise involving people can be successful, let alone competitive and sustainable, if no attention is given to fostering the knowledge life-cycle. I would argue that that is true whether the enterprise is or is not for-profit. All organizations face competition, unless they are a monopoly. Not-for-profit organizations face competition from other service providers, as well as from the possibility the individual will choose not to use the services. This


* Paper presented at the Special Libraries Association 93rd Annual Conference Los Angeles California USA June 8 - 13, 2002

paper, however, will only discuss knowledge management in the context of a for-profit enterprise.

The business case for formally managing the knowledge life-cycle (i.e. knowledge management) has gotten easier to make, for all types of firms, as the understanding of core competencies has grown. It is generally accepted today that a firm competes and gains competitive advantage based on its core competencies. Core competencies represent the firm's capacity to create new businesses by creatively combining core skills with new knowledge to form new "knowledge streams".¹ For example, chemical firms need to blend chemistry and chemical engineering with software engineering and electronics so as to automate production operations, while consulting firms must effectively combine classic management consulting with the delivery of cutting edge information technology solutions. Competitive advantage is gained by the ability to become competent in doing this in an environment of dynamic change. To become competent, the organization has to be able to repeatedly learn and transfer the learning to new team members. The ability of the competence to be shared, therefore, is as critical as the other components of a true core competency. A firm must be agile in developing and repurposing core competencies as new business opportunities present themselves. This requires the ability to transfer tacit and explicit knowledge (people-embodied knowledge), and skills (capital-embodied knowledge). "It is the combination of people-embodied and capital-embodied knowledge that represents the totality of the competence base within an organization." (Prahalad p. 241) Classic examples of core competencies include miniaturization at Sony, network management at AT&T, chemical engineering at DuPont, and user-friendliness at Apple.

The business case for knowledge management, therefore, rests entirely on the assertion that purposefully managing the knowledge life-cycle within an enterprise is the critical enabler in fostering the development of new core competencies as well as sustaining existing ones. Without core competencies, a firm has no ability to differentiate itself in the marketplace and will not be viable. For simplicity, this paper will only address the narrow portion of KM that is the management of knowledge assets (i.e. the acquisition and handling of explicit knowledge).

¹ Prahalad, C. K.; Fahey, Liam, and Randall, Robert M. *Creating and Leveraging Core Competencies*. Fahey, Liam and Randall, Robert M., Editors. *The Portable MBA in Strategy*. 2nd ed. New York: John Wiley & Sons, Inc.; 2001; pp. 236-252.




Context Matters

While all enterprises must have core competencies to successfully compete in the marketplace and all enterprises must formally manage their knowledge in order to support the creation and maintenance of those core competencies, there are subtle, yet distinct, differences in the resulting knowledge assets. These differences result in different requirements with regard to standards, methods, and processes for the management of the assets. For example, what makes a consulting firm different from a chemical manufacturing firm? Without going through the tedium of analyzing the different business models, the difference ultimately is embodied in the difference of the primary knowledge base supporting their core competencies. As Prahalad aptly emphasized, core competencies are derived from the combination of people-embodied and capital-embodied knowledge available to the enterprise. I would argue that the different industries within which firms compete define the viable business models available to the firms, and those business models express the constraints which drive the differences between how knowledge assets are managed in a manufacturing company versus a consulting firm. Not surprising, it is ultimately the distinctly different business drivers that define the look, feel and operability of the knowledge asset management system; form follows function. After all, knowledge management must be targeted to meet business needs, which requires it reflect the manner in which the business operates.

People-embodied knowledge includes tacit and explicit knowledge. Tacit is the elusive, not yet articulated wisdom and knowledge we all have, which has been shaped by our experiences, learning, understanding and intellectual capabilities.

Explicit knowledge is articulated and codified understanding in a form capable of being shared and transferred. The life-blood of consulting firms is the exchange between tacit and explicit knowledge. Capital-embodied knowledge is explicit knowledge which may be either internally generated or purchased from external resources. It is a combination of knowledge and skills made explicit to only those within the enterprise (proprietary), and purchased knowledge and skills from external experts. Chemical manufacturing companies rely heavily on capital-embodied knowledge. Patents, trade secrets, technical skills, research & development and strategic partnerships are all critical to the sustainable competitive advantages of chemical companies. So while the chemical manufacturer and the consulting firm each have some degree of business reliance on both people-embodied and capital-embodied knowledge, value-creation by the enterprise depends differently on each type of knowledge.



In addition to the differences in the extent of dependence on the two types of knowledge, different types of firms apply knowledge at different points in the process of value creation. Michael Porter's value chain model maps the creation of value by the firm, starting with suppliers and ending with customers. A typical chemical manufacturing firm creates value as a result of its relationships with suppliers, its raw material procurement efficiencies, the depth and breadth of its R&D (future product pipelines and support of existing product lines), and its manufacturing competencies. On the other hand, and at the other extreme, a consulting firm creates value primarily as a result of the quality of its relationships with its customers. Those relationships are highly dependent upon the firm's marketing skills (proposal development), the ease of service distribution via local offices located globally, and the overall competitive quality of the delivered services. The two types of firms create value at different ends of the value chain. As a result, the firms rely on different types of knowledge to different degrees (people-embodied vs. capital-embodied), and they create different types of knowledge assets. As a result, the firms have very different system requirements for the creation, capture, storage, retrieval, use and repurposing of those assets; form follows function.

In theory, there are standards

Any building can house thousands of books and journals, and any computer can house gigabytes of digitized content, but it is the core competencies of the information science profession that enable users to make sense out of the jumble of information and data. In an ideal world, it is the core competencies of the information science profession that enable users to precisely retrieve the right knowledge for the competitive advantage of the business.

Information Science Core Competencies: Sense-Making Tools

1. Expert knowledge of information resources (content and use)
2. Conceptual analysis (indexing, abstracting)
3. Ability to structure and organize content (information management)
4. Ability to synthesize and customize (information relevancy)

We all learned in library school about the standards and guidelines of collection management, cataloging and classification schemes, as well as the creation and use of thesauri and authority lists. Whether the knowledge assets are in microform, print or digital format, there are standardized work processes to be established and followed with regards to collection management and classification. The intent of the standards is to ensure *sustainability and consistency* of access to the knowledge assets *over time and across disciplines*. The standards are intended to ensure the right content is collected, at the right time, in the right format and stored in a manner facilitating fast and accurate retrieval at the right time. Striking a balance between relevancy and recall during the retrieval process is the ultimate goal. It is as equally damaging for an information seeker to miss critical content because the retrieval set is too narrow, as it is for them to lose precious time sifting through too much content because the retrieval set is too broad. As important as that balance is in an academic or public library, it is that much more important in a corporate setting, where *speed to market* is critical for survival.

For the sake of brevity, we will look at just two of the standardized information processes; content management and content classification. Content management includes defining the scope of the collection, establishing the acquisition guidelines, identifying storage and retrieval mechanisms, and determining retention policies.

- Step 1: Define the Scope of the Collection.
 - Based on business need, what are the primary and secondary subject areas?
 - Do certain topics require greater depth than others?
 - Do certain topics require greater breadth than others?
 - What are the characteristics that make objects within those topic areas attractive to collect? In other words, how is quality defined and assessed?

- Step 2: Establish Content Acquisition Guidelines
 - Where is the content generated, in what format, and by whom?
 - Will content be purchased, generated internally or both?
 - How will the selected materials be obtained? What processes are required to support the acquisition?

- Step 3: Create Storage and Retrieval Mechanisms
 - Based on expected user access points, where will the collection be stored? Electronic repository, physical library, a combination?
 - What are the various formats for the content? All digital? All print? All microform? A combination?
 - What will the search and retrieval mechanisms be? How will the content be catalogued or indexed? What processes and technologies are required to support storage and retrieval?
 - How will end-users be educated to effectively use the tools?

- Step 4: Determine the Retention Policies
 - Does retention vary by content type? By topic? By author?
 - Do different topic areas have different currency life-spans?
 - If the material is removed from the active collection, is it kept accessible in an “archive” collection?

All of these aspects of content management are relatively straightforward when the collection is associated with an established and relatively static discipline or topic area. Keeping with our comparison, if we consider the collection within a chemical manufacturing firm, we see that the standards and guidelines easily

apply with minimal adjustment to business needs. In fact, the print/microform collection at a chemical manufacturing firm looks and operates much the same as a library at a university. The print collection is catalogued using Library of Congress classification scheme. The OPAC follows standard MARC record format. Proprietary research literature is indexed and abstracted. The abstracts are stored in a citation database and searchable at every employee's desktop using Boolean logic. The scope of the collection is defined by the nature of the business: polymer chemistry, chemical engineering, physics, etc. Some sub-disciplines may be more important to the firm, and so may have greater depth in the collection. The knowledge assets in the collection are both internally generated and purchased from external vendors. Retention policies vary based on content type. Proprietary content is permanent, as are the majority of the purchased monographs and reference materials. Periodicals have various retention periods, as determined by the currency of their content. Peer-reviewed technical journals are kept permanently if space is available. Newsletters and trade journals are kept as short as 3 months and up to 24 months as defined by frequency of publication and currency of content. As will be seen later, none of the four steps enumerated above are easily applied "as is" to knowledge asset collections in a consulting firm.

The second standardized process considered here is the classification of content and the subsequent cataloguing or indexing of the material. Conceptual analysis, one of the core competencies of information science, is the process of determining the primary topic(s) embodied in a knowledge asset and then applying the appropriate Library of Congress call number, controlled term, or subject heading. The processes of cataloging and indexing are intended to ensure the content is stored in a manner facilitating fast and accurate retrieval when the material is needed. The use of standardized classification schemes (e.g. LOC), controlled terms from a thesaurus (built to ANSI/NISO Z39.19-199x specifications), or standardized subject headings (e.g. MESH) ensures consistent ability to precisely and comprehensively retrieve content across repositories. The consistent application of classification standards across repositories enables information seekers to apply existing understanding of the scheme to efficiently retrieve content across time and disciplines. Without standards, information seekers would be lost in the unique processes required to retrieve content from each different repository.

Classification standards are relatively straightforward to apply to the collections of knowledge assets at the chemical manufacturing firm. The various branch libraries all use LOC call numbers to catalogue monographs and reference collections. The OPAC, using MARC format, contains LOC subject headings and

adheres to AACR. The proprietary document collection is indexed using controlled terms from an evergreen thesaurus. Because the volume of new knowledge assets added to the collection is comparatively low, the conceptual analysis of the documents is efficiently done by information professionals with appropriate scientific domain knowledge. As will be seen shortly, significant compromises must be made in the consulting firm with respect to classification standards and conceptual analysis processes.

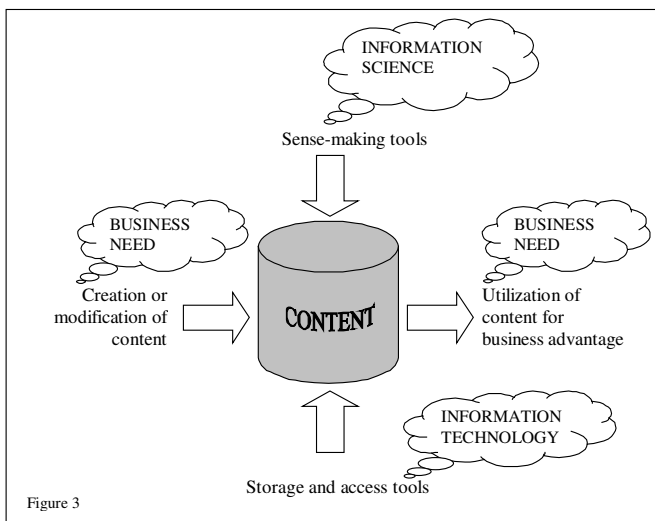
In practice, there are compromises

We expect businesses to have an appreciation for the sense-making tools of information science, but the fact is that they seldom do. To be heard, we need to do more than just assert that our tools transform jumbles of information into usable repositories of knowledge assets. Information scientists must have an appreciation for when standards are called for and when compromises are called for. The first step in generating the business's appreciation for information science core competencies is to demonstrate an understanding of the needs of the business. Applying standards for the sake of applying standards, with no adjustment for business needs, will drastically reduce the ability of the knowledge management system to provide business value. Corporate knowledge bases do not exist just to exist. They exist to meet business needs, and therefore, must be adjusted to ensure reliability, relevance and business value. In a consulting firm environment, this means many compromises are called for.

Despite having extensive experience applying information science standards and guidelines in a corporate environment, to both a digital knowledge base of proprietary research and a multitude of special print & microform collections of published literature, the reality of doing so in a consulting firm environment has been one long lesson in compromising. The chronic, yet welcome, challenge is to find the balance between the standards which will ensure consistent retrieval of content across time and disciplines, and the needs of the business. How to adjust form to follow function? The standardized process and guidelines that work well in a comparatively slow moving business environment (chemical manufacturing) simply can not be applied in their native format in a fast paced consulting environment. The challenge is in finding a way to apply them so that business processes are supported and not hindered.

As in any corporation, the intent of a consulting firm's knowledge asset repository is to capture the explicit firm knowledge generated during the course of business operations and make it available to support additional revenue generation. The assets are stored in the repository, where information science's "sense-making"

tools are applied so as to ensure the content can be retrieved and utilized for the competitive advantage of the business. (Figure 3) In most corporate environments, the standards and guidelines associated with content management and classification can be applied to the knowledge asset repositories with little or no adjustments. For example, a polymer chemist writes a research report discussing their work on a research program, and contributes the document to the repository. The information scientist with polymer chemistry domain knowledge scans the document, indexes it appropriately with controlled terms from the thesaurus, writes an abstract, submits the citation with indexing and abstract to the citation database and finally sends the report to the physical central repository. The indexing will have captured the essence of the content. It will not take into account how someone might utilize the content, what business implications might be inferred, or anything else. As per standardized cataloguing rules, the indexing reflects precisely what is inherent in the document: standard bibliographic citation fields plus the major concepts. In this case, this is primarily because knowledge assets in the manufacturing firm are frequently re-used but rarely repurposed. They are used in the same manner they are created. This polymer chemistry research report will be pulled and reread when it is time to build on the research discussed in the document.



In a consulting firm, however, knowledge assets are repurposed far more often than reused in their native form. Most practitioners take components from various assets in order to synthesize new ideas, which are then captured within new knowledge assets. The knowledge asset is rarely valuable in its entirety “as is”. Rather, its value stems from the many ways the knowledge in the asset can be repurposed. Portions of the material can be extracted and used in a marketing piece (“we did this sort of work for so-and-so”). Other portions can be combined with bits from other assets to capture the firm’s thinking on a specific business issue. There are an endless number of possibilities. In addition, the content generated by one practice area of the consultancy may be of direct relevance and import to one or more other practice areas. However, what creates that value may vary between the practice areas. For example, a sample piece of content may address strategic redesign of the supply chain and implementation of an IT-based solution for a client operating in the aerospace & defense industry. There are a minimum of at least three practice areas (disciplines) to which this content would have direct relevance: strategy, supply chain, and aerospace & defense. Each practice area has “ownership” of the content, has their own spin on it (resulting in different abstracts) and has different retention policies for this single piece of content. The content may be critical to a new service offering in one practice area, and old hat in another. Also, each practice area may repurpose the material in one of three different ways over time. It could be used for project support, research, or sales & marketing. The point is that the knowledge management team for each of the different practice areas could be applying different cataloging and indexing rules, following different retention policies, and writing different abstracts all for the same knowledge asset with a single database record in the central repository. Each team will emphasize the aspect of the content which has the greatest relevance to their constituency group. In the process of entering the asset into the collection, the knowledge managers appear to have compromised on nearly every aspect of collection management and classification standards and guidelines. The standards and guidelines are adjusted to reflect the nature of the content for the practice area, the use of that content, the extent of churn, the level of participation from the practitioners, the level of user sophistication, etc.

COMPROMISE 1: Manage knowledge where it’s created, but do that within a global system.

Stewart emphasized that knowledge must be managed within the context where value is created.² In a consulting firm, knowledge and value are created

² Stewart, Thomas A.; *The Wealth of Knowledge*. New York: Currency; 2001; ISBN: 0-385-50071-8, p 180.

at the practice area level. Yet, if the firm's knowledge was literally managed at the practice area level, the firm would be at a tremendous disadvantage in the marketplace. There would be multiple silos of knowledge in the firm with no cross-fertilization between practice areas. One way to avoid silos is to have a single firm-wide repository for the knowledge assets. This first compromise then triggers the subsequent compromises because of the differing business requirements of the various practice areas within the firm.

COMPROMISE 2: No single scope defined for the mega-collection within the repository, but clear scopes defined within practice area collections.

The scope of the collection (content type and topic) is an aggregated scope determined by the business needs of the various specific practice areas. Acquisition guidelines and processes are established based on consideration of the knowledge asset creation rate, the workload of practitioners, their ability to access the firm's intranet from remote locations, etc. All of these aspects differ substantially across the practice areas. One area may focus on collecting only proposals while another may focus primarily on collecting completed project deliverables. One collection is focused on supporting sales efforts while the other is focused on supporting service delivery.

COMPROMISE 3: Inconsistently index content across the entire repository but consistently within a "practice area".

Storage is dictated by the firm requirement that ALL knowledge assets are to be stored in a single firm-wide repository, and so, to some degree the retrieval mechanism, is also predetermined. Practitioners can either use the full-text search engine, or use the handful of captured attributes of the content to browse their way to the knowledge assets. These attributes include client, client industry, document type, geographic region, etc. While conceptual analysis for concept indexing is not done across all practice areas, it is done within some of the areas. For example, the manufacturing practice area periodically generate lists of controlled terms, no more than 5 to 6 "topics of lasting value", and assign those keywords against the content. This additional, consistently applied indexing significantly enhances the precision of content retrieval for the practitioners.

COMPROMISE 4: Follow the money, regardless of the 'true' nature of the asset.

One would expect that at least there would be no need for compromises with respect to the core attributes of the knowledge asset. However, compromises are made here too. One attribute used to index the content is the client for whom the work is done. All clients operate in a primary industry. Yet, the

client may not be indexed by their true primary industry. They are in fact indexed by which industry practice is credited with the revenue from that client. For example, a large pharmaceuticals firm may be indexed with the chemical manufacturing industry because the *revenue* from the client is realized by the manufacturing practice area, even though a much more appropriate tag to assign would be ‘pharmaceutical manufacturing firm’ in the ‘healthcare industry’.

COMPROMISE 5: The sky is not always blue, even when it is blue.

Cataloguing and indexing standards require that nothing be inferred about the item being classified. If the project work was done in the manufacturing industry, then it is indexed with ‘manufacturing’. It is left to the end-users to have enough domain knowledge to understand that there may be content of equal interest and applicability to their query in the consumer business or health care industries. However, in the consulting firm’s repository, this domain knowledge is not presumed. Content is indexed based on not only its inherent characteristics but also based on what might be inferred about the material. So, while the client work was done for a manufacturing client, the industry tags would include all related industries.

COMPROMISE 6: Retain dated content because it is fresh content for other repository constituencies.

Finally, as could be predicted, the retention policies of one practice area often come into conflict with those from another practice area. The content type or the business issue addressed or the client for whom the work was done could be insignificant for the industry practice area, but for the involved competency or service practice area, it may represent either their core competency, a best practice or be cutting edge service delivery. The industry practice area would not want to keep the content long, while the service area may want to keep it for years. The industry practitioners get frustrated because this “bad” hit continues to be returned in their query answer sets, and if it were archived, then the competency practice area would be missing a vital knowledge asset.

Through a consensus process, chronic trade-offs between aspects of information architecture and design, information storage and retrieval, collection development and management, standards and guidelines, and virtual community development are settled upon by the knowledge management teams within practice areas and across the firm. The result is a form of contained chaos that succeeds at supporting business objectives across the firm. In fact, the management of knowledge assets models how the practice areas create, use and repurpose

knowledge; form follows function. Almost all client service projects are bid on, won, and completed by multi-practice area teams, with each practice area delivering its own twist to the work.

Conclusions

I no longer believe that compromising information science standards and guidelines is by definition bad implementation of information science. The success of a knowledge management program in a for-profit setting is determined by the extent to which the program supports and even enhances the firm's ability to go to market (i.e. generate revenue). Success is not determined by the extent to which the system adheres to information science standards and guidelines. Those standards and guidelines have value only to the extent they do not hinder the firm's operations while they ensure the sustainability of the knowledge management system. I would challenge, in fact, that the six compromises noted above are not really even compromises. Instead they represent an iterative process of determining a finer level of granularity against which the standards can be consistently *applied*. Perhaps the only compromises are:

- first recognizing that what we learned in graduate school needs to be applied at the molecular level rather than the level of the organism;
- and second, working across practice area boundaries to find the lowest common denominator against which to apply standards.

The lowest common denominator is not the knowledge asset, but the knowledge captured in that asset. The bits of knowledge embodied in the document are the molecules; the document is the organism.

At first blush, it looks as if there is no information science standards employed in the knowledge asset repository, when in fact there are multiple sets of standards cohabitating. Driving to the lowest common denominator of base standardization leads to the minimization of costs and maximizes flexibility. For example, the volatility of terminology in the consulting industry and the speed at which the business focus can change makes the indexing of knowledge assets in the manner which a chemical firm can, economically infeasible. Nor is there much of a business case for trying to. Unlike the terminology in scientific disciplines, consulting terminology is both a communications vehicle used to convey concepts and, probably more importantly, a sales and marketing tool. As a result, consulting language has a high churn rate. Supply chain management became value chain management which became vertical management which became collaborative commerce. The cost of re-indexing the material tagged with the

older concept would be prohibitive. Even managing a thesaurus that would educate the search engine as to the equivalency of those terms becomes prohibitive, especially when you look at the alternative of simply relying on free-text searching, the end-users' ability to use their domain knowledge to search effectively, and the relatively short life-span of a knowledge asset in a consulting firm (18 months on average), why bother spending the time and money? Consulting firms can not afford to use resources to re-index content and restructure keywords and subject headings when business service offerings and focus change as rapidly as they do in the consulting industry.

The standards and guidelines of information science, therefore, are not really compromised, but the core competencies of information science are challenged to truly demonstrate their value and adaptability. The successful application of information science standards and guidelines to ensure sustainability and consistency of such dynamic and evergreen knowledge bases in fact demonstrates the tremendous business value of such competencies. By implementing information science standards at the level of the lowest common denominator, the rich potential of the knowledge repository is significantly deepened. For example, the finer level of granularity allows the flexibility to selectively repackage individual pieces of content into hyperlinked collections of content that address a current "hot-topic" without disrupting any existing classification or indexing against the various component knowledge assets of this new virtual special collection. The knowledge assets in the virtual collection were not previously related to each other and the business value of the collection may be short lived. As critical focus areas for the firm change, content can be rapidly disassembled and reassembled as needed to assist practitioners in rapidly delivering results to the firm's clients.

Finally, form does follow function in a successful knowledge management system. The level of structure and the degree of constraints within the system mirror the business supported by the knowledge asset repository. The level of granularity against which the standards are applied is dependent upon the nature of the business.

Martha K Heyman
15 Randy Lane
Cochranville, PA
USA 19330
mheyman@dc.com